

The benefit of a 1-bit jump-start, and the necessity of stochastic encoding, in jamming channels

Bikash Kumar Dey, Sidharth Jaggi, Michael Langberg, Anand D. Sarwate

February 9, 2016

Abstract

We consider the problem of communicating a message m in the presence of a malicious jamming adversary (Calvin), who can erase an arbitrary set of up to pn bits, out of n transmitted bits $\mathbf{X} = (x_1, \dots, x_n)$. The capacity of such a channel when Calvin is *exactly causal*, *i.e.* Calvin's decision of whether or not to erase bit x_i depends on his observations (x_1, \dots, x_i) was recently characterized [1, 2] to be $1 - 2p$. In this work we show two (perhaps) surprising phenomena. Firstly, we demonstrate via a novel code construction that if Calvin is *delayed* by even a single bit, *i.e.* Calvin's decision of whether or not to erase bit x_i depends only on (x_1, \dots, x_{i-1}) (and is independent of the “current bit” x_i) then the capacity increases to $1 - p$ when the encoder is allowed to be stochastic. Secondly, we show via a novel jamming strategy for Calvin that, in the single-bit-delay setting, if the encoding is deterministic (*i.e.* the transmitted codeword \mathbf{X} is a deterministic function of the message m) then no rate asymptotically larger than $1 - 2p$ is possible with vanishing probability of error; hence *stochastic encoding* (using private randomness at the encoder) is essential to achieve the capacity of $1 - p$ against a one-bit-delayed Calvin.

1 Introduction

There are two traditional methods in information theory for modeling uncertainty in communication channels. Shannon's approach treats uncertainty in the channel as a random phenomenon and requires the probability of decoding error to vanish as the blocklength tends to infinity [3]. The capacity is governed by the behaviour of typical channel realizations; for example in a binary erasure channel (BEC) with erasure probability p , the channel will erase *approximately* pn symbols as $n \rightarrow \infty$. Classical error-control coding, which we might call Hamming's approach, considers the problem of worst-case recovery. Assuming the channel erases *at most* pn symbols, the goal is to design codes that can exactly recover the transmitted message.

One way to view the differences between these two models is to anthropomorphize the channel and assume it is being controlled by an adversary (whom we call Calvin), who wishes to foil the communication between the transmitter and receiver (hereafter referred to as Alice and Bob). By restricting the information available to Calvin we can recover models for communication in these two regimes. This information could be about the transmitted message or the codeword itself. For example, a BEC could be modeled by an *oblivious* adversary who knows neither the message nor the codebook used by the transmitter and receiver, and is restricted to erase no more than pn symbols as $n \rightarrow \infty$. In the BEC we allow for some probability of error (average or maximum over messages) that tends to 0 as $n \rightarrow \infty$. The Hamming approach is more pessimistic: Calvin knows the transmitted message, codeword, and codebook, and can adversarially choose up to pn positions




Figure 1: Binary adversarial erasure channels.

to erase to create uncertainty at the decoder. A good code in the Hamming sense protects against all such erasure attacks and guarantees zero error subject to the adversary's constraint.

The advantage of this (perhaps paranoid) adversarial modeling is that it reveals a plethora of intermediate models between the Shannon and Hamming models that can potentially shed light on the difference between average and worst case analysis. An arbitrarily varying channel (AVC) [4] has a time-varying state (e.g. the presence/absence of an erasure) that can be chosen by Calvin. In AVC models, distinctions between error criteria (maximum or average) and the presence of common randomness shared by Alice and Bob become important [5]. Sometimes the AVC capacity displays a dichotomy: if Calvin can simulate sending a legitimate message, then Bob may not be able to decode correctly. Such AVCs are called *symmetrizable*, and the capacity is 0 in this case [6].

In this paper we find a new dichotomy when Calvin can observe the transmitted codeword subject to some *delay*. That is, at time i , Calvin has knowledge of the transmitted codeword up to time $i - \Delta$. In particular, we study the case $\Delta = 1$ for a model in which Calvin can erase at most pn of the transmitted bits. If the encoder and decoder share common randomness, then prior work shows that the capacity in this model is $1 - p$, the same as the BEC capacity [7].

We show that in our model of study with $\Delta = 1$ the capacity is $1 - p$. Our coding scheme uses randomness at the encoder. Specifically, for any rate below $1 - p$, the maximum probability of error (over the encoder randomness) goes to 0 as $n \rightarrow \infty$. This result may come as a surprise, as the capacity for $\Delta = 0$ is strictly lower and equals $1 - 2p$ [1, 2]. Moreover, we show encoder randomness is essential by proving that any deterministic coding scheme will have capacity at most $1 - 2p$ for $\Delta = 1$. In contrast, we show in Sec. 5 that for omniscient adversary, who has noncausal knowledge of the full codeword, the capacity under stochastic encoding is the same as that under deterministic encoding.

1.1 Prior work and contributions

We focus on two aspects of communication models with adversaries: the impact of delay on the knowledge of the adversary, and the difference between deterministic and stochastic encoding. The first paper to our knowledge that examined these issues was by Ahlswede and Wolfowitz [8], who gave several equivalences between classes of AVC models and further showed that stochastic encoding alone can have some benefit over deterministic encoding. Traditional works on the AVC [5]

focused on the case where Calvin is oblivious ($\Delta = n$); for erasure adversaries the capacity for average error and deterministic codes is $1 - p$ [6]. If the encoder and decoder share common randomness then the capacity is $1 - p$ even if Calvin is omniscient ($\Delta = -n$) [9]. For deterministic codes the best-known achievable rate equals $1 - H(p)$ via GV codes [10, 11] and the best-known outer bound is given by the LP bound [12].

Our results show a sharp difference between $\Delta = 0$ and $\Delta = 1$. For $\Delta = 0$, Bassily and Smith proved an outer bound of $1 - 2p$ and Chen et al. [2] constructed a code with stochastic encoding that achieves $1 - 2p$. For $\Delta = -\epsilon n$ (that is, ϵn -lookahead) this code achieves $1 - 2p - \epsilon$, showing that sublinear lookahead cannot improve Calvin’s jamming strategy. We find the capacity for $\Delta = 1$ (hence the title “one-bit delay”) is $1 - p$, thereby establishing the same result for all positive Δ . This demonstrates a very sharp asymmetry between the effect of lookahead and delay! The capacities of the problems in the spectrum above are plotted in Figure 1.

The second issue we address is the importance of stochastic encoding (private randomization). The first paper on AVCs [4] considered the case $\Delta = 1$ with full common randomness, but their proof does not extend to constrained adversaries [13]. Most AVC results focus on the difference between common randomness and deterministic coding for oblivious [4, 14, 13, 6] or omniscient [15, 16, 9] adversaries. Stochastic encoding offers few benefits in these settings for DMCs or AVCs, although it is useful in wiretap scenarios [17]. In this paper we show that for $\Delta = 1$ deterministic codes cannot achieve rates higher than $1 - 2p$ whereas stochastic encoding can achieve a rate $1 - p$. This shows that stochastic encoding is essential for the specific channel considered in this work. A related (and fascinating) open question is whether the same is true for exactly causal binary erasure channels – the rate-optimal codes [2] achieving $1 - 2p$ used stochastic encoding, and it is unclear whether the same rate is achievable via deterministic codes. However, for an omniscient adversary, the capacity under stochastic encoding is argued to be the same as that under deterministic encoding in Sec. 5.

1.2 For comparison: Large alphabet channels

Often, analyzing “large alphabet” channels (where the channel input/output alphabet sizes are larger than the blocklength n) gives one insight about general channels, including channel models that are challenging to characterize (such as binary channels).

For large alphabet erasure channels, the situation is somewhat different than considered in this paper. If the input alphabet \mathcal{X} is of size q which is at least n , and at most a p -fraction of symbols may be erased, the capacity is exactly $1 - p$. (This equals the capacity of the q -ary random erasure channel, in which the erasure probability of each symbol is p .) This rate is attainable regardless of the knowledge of the adversary, and computationally-efficiently attainable by (deterministic) Reed-Solomon codes; hence neither of the behaviours observed in the binary adversarial erasure channel is observed here.

On the other hand, for large alphabet symbol errors when at most pn output symbols may differ from the input symbols, we may observe similar behaviour to the binary erasure case. In [18], it was demonstrated that the capacity of exactly causal channels equals $1 - 2p$, which is the same as the capacity if the adversary is omniscient – hence there is no advantage of lookahead for the adversary. In both cases we can achieve these rates using computationally efficient (and deterministic) Reed-Solomon codes. However, if the adversary is delayed, then, depending on the symbol-error model, the capacity may be higher. Two symbol-error models were considered. When symbol errors are *additive* (the output symbol y_i equals $x_i + e_i$, where x_i is the input symbol, e_i is the error symbol, at most pn e_i s may be non-zero, and all symbols and addition are over the finite field \mathbf{F}_q) with a delay

of even a single symbol, the capacity equals $1 - p$ (thereby exhibiting a similar phase-transition in the capacity as in this paper). In contrast, with *overwrite* errors (the output symbol y_i equals e_i for at most pn non-zero e_i s) with a delay of dn symbols (e_i can be a function of x_1, \dots, x_{i-dn}), the capacity is $1 - 2p + d$ for $p < 1/2$, and 0 otherwise, hence demonstrating a less sharp transition in the throughput.

The differences in optimal rates obtainable with stochastic and deterministic encoding over large alphabets with causally-constrained adversaries has not, to the best of our knowledge, been considered in the literature and may be worthy of investigation.

2 Channel model

For integers $r < s$ let $[r : s]$ denote the set $\{r, r+1, \dots, s\}$ and let $[N]$ denote the set $[1 : N]$. For a set $S \subseteq [n]$, let \bar{S} be the complement of S . Let \perp denote the erasure output symbol. Random variables will typically be denoted by capital letters and vectors by boldface. For a vector $\mathbf{z} = (z_1, z_2, \dots, z_n)$ and set $S \subseteq [n]$ we will write \mathbf{z}_S for the vector $(z_i)_{i \in S}$ with the components ordered in increasing order of index. The Hamming weight of a binary vector \mathbf{z} is $\text{wt}_H(\mathbf{z})$, and Hamming distance is d_H .

We first set up our channel model more generally before specializing to the case considered in this writeup.

Let \mathcal{X}, \mathcal{Y} , and \mathcal{Z} be discrete alphabets. We consider variants on arbitrarily varying channel models, which are channels whose state $z \in \mathcal{Z}$ is (partially) controlled by a malicious adversary who wishes to prevent reliable communication across the channel. The model is parameterized by a set of discrete channels $\{W(y|x, z) : x \in \mathcal{X}, y \in \mathcal{Y}, z \in \mathcal{Z}\}$. For blocklength n , input $\mathbf{x} \in \mathcal{X}^n$, state $\mathbf{z} \in \mathcal{Z}^n$, and output $\mathbf{y} \in \mathcal{Y}^n$, the blocklength- n extension of this channel is

$$W(\mathbf{y}|\mathbf{x}, \mathbf{z}) = \prod_{i=1}^n W(y_i|x_i, z_i). \quad (1)$$

An $(n, 2^{nR})$ code with randomized encoding for this channel is a pair of maps (Φ, ψ) where $\Phi : [2^{nR}] \rightarrow \mathcal{X}^n$ is a randomized encoding map and $\psi : \mathcal{Y}^n \rightarrow \{0\} \cup [2^{nR}]$ is a deterministic decoding map. In a deterministic code, the encoder is also deterministic, and it assigns a unique codeword to each of the 2^{nR} messages.

We consider channel models in which \mathbf{z} is chosen adversarially and with partial knowledge of the transmitted codeword. We define an adversarial strategy Γ of delay Δ to be a sequence of maps $\{\gamma_t : t \in [n]\}$, where $\gamma_t : \mathcal{X}^{t-\Delta} \rightarrow \mathcal{Z}$ is a randomized map from $\mathbf{x}_{[1:(t-\Delta)]}$ to z_t . We allow this map to depend on the code (Φ, ψ) . Alternatively, such a strategy defines a conditional probability distribution $G(z_t|\mathbf{x}_{[1:(t-\Delta)]}, \Phi, \psi)$ which chooses Z_t . The corresponds to a scenario where the adversary can observe the channel input up to delay Δ and can choose the channel state based on that information and the structure of the code. We say the strategy satisfies a cost constraint p with respect to the cost function $c : \mathcal{Z} \rightarrow \mathbb{R}^+$ if

$$\sum_{t=1}^n c(Z_t) \leq pn. \quad (2)$$

Let $\mathcal{G}(p, \Phi, \psi)$ be the set of strategies that satisfies the cost constraint.

The probability of error for this code on message $m \in [2^{nR}]$ with adversarial strategy Γ is

$$P_{\text{err}}(m, \Gamma) = \sum_{\mathbf{y}: \psi(\mathbf{y}) \neq m} \sum_{\mathbf{x} \in \mathcal{X}^n} \sum_{\mathbf{z} \in \mathcal{Z}^n} W(\mathbf{y}|\mathbf{x}, \mathbf{z}) \mathbb{P}(\Phi(m) = \mathbf{x}) \prod_{t=1}^n G(z_t | \mathbf{x}_{[1:(t-\Delta)]}, \Phi, \psi). \quad (3)$$

The maximum probability of error is

$$P_{\text{err}} = \max_{\Gamma \in \mathcal{G}(p, \Phi, \psi)} \max_{m \in [2^{nR}]} P_{\text{err}}(m, \Gamma). \quad (4)$$

Note that these probabilities are over the encoder randomness in Φ , potential randomness in the adversary strategy Γ , and possible randomness in the channel. We say a rate R is achievable in this model if there exists a sequence of $(n, 2^{\lfloor nR \rfloor})$ codes such that $P_{\text{err}} \rightarrow 0$ as $n \rightarrow \infty$. The capacity is the supremum of the set of achievable rates.

Here we take $\mathcal{X} = \{0, 1\}$, $\mathcal{Z} = \{0, 1\}$, and $\mathcal{Y} = \{0, 1, \perp\}$. The channel model is given by $y_t = x_t$ if $z_t = 0$ and $y_t = \perp$ if $z_t = 1$. The cost function is $c(z) = z$ and $\Delta = 1$. This corresponds to a binary-input channel in which the adversary can observe all past inputs and can erase up to pn of the bits. Under a larger delay the capacity is $1 - p$. However, for delay 0 the capacity is $1 - 2p$. In the remainder of the paper we will show that with stochastic encoding the capacity is $1 - p$ and with deterministic encoding the capacity is at most $1 - 2p$.

Our main results take the form of two theorems. Theorem 1 is an achievability result: it says that the stochastic encoding can achieve rate $1 - p$ against a bit-erasing adversary who can erase up to pn bits and is subject to delay $\Delta = 1$.

Theorem 1. *The capacity of a binary channel with a bit-erasing adversary who can erase up to p fraction of a codeword based on causal 1-bit-delayed observation is $1 - p$.*

The next theorem contrasts the above result to say that if the transmitter is restricted to using deterministic codes, then the capacity is at most $1 - 2p$. Therefore stochastic encoding is crucial to take advantage of the adversary's delayed observation.

Theorem 2. *The capacity of a binary channel under deterministic encoding, with a bit-erasing adversary who can erase upto p fraction of a codeword based on causal 1-bit-delayed observation is at most $1 - 2p$.*

3 Analysis for Stochastic Encoding: Proof of Theorem 1

We consider coding for an online adversarial channel with binary inputs in which the adversary observes the channel input subject to unit delay and can erase a fraction p of the bits. Under a larger delay the capacity is $1 - p$ bits. However, for delay 0 the capacity is $1 - 2p$ bits.

3.1 Code construction, encoding, and decoding

Given a parameter $\epsilon > 0$, rate $R = 1 - p - \epsilon$ and blocklength n , let $M = \lfloor 2^{nR} \rfloor$ be the number of messages.

3.1.1 Random code construction

Our code construction relies on the following parameter settings:

$$K = (1/4) \log_2 n \quad (5)$$

$$q_k = 2^{k-1} n^{-1/2}, \quad k \in [K] \quad (6)$$

1. For each message $m \in [M]$ there is a *base codeword* $\mathbf{u}(m)$, selected uniformly at random from $\{0, 1\}^n$.
2. For each message $m \in [M]$ the encoder has a partition $\{S(m, k) : k \in [K]\}$ of the set $[n]$ so that for each $k \in [K]$ the set $S(m, k)$ is a set of indices of the codeword. We generate the partitions $\{S(m, k)\}_k$ for each m by binning the indices in $[n]$ into K bins independently and uniformly at random.
3. In addition the encoder maintains a set $\mathcal{Q} = \{q_k : k \in [K]\}$ of probabilities, where K and q_k are given by (5) and (6).

3.1.2 Encoding

The encoding is randomized. To encode the message $m \in [M]$ the encoder transmits $\mathbf{X} = \mathbf{u}(m) \oplus \mathbf{Z}$, where $\mathbf{Z} = (Z_1, Z_2, \dots, Z_n)$ with $Z_i \sim \text{Bernoulli}(q_k)$ if $i \in S(m, k)$. That is, for each $k \in [K]$, the encoder adds a $\text{Bernoulli}(q_k)$ noise to the components indexed by $S(m, k)$. Our encoder is depicted in Figure 2.

3.1.3 Decoding

1. Given the received codeword \mathbf{Y} , the decoder first finds the smallest index τ such that the first τ positions of \mathbf{Y} contain $(R + \epsilon/2)n$ unerased bits:

$$\tau = \min\{t : |\{i : i \leq t, \mathbf{Y}_i \neq \perp\}| \geq (R + \epsilon/2)n\}. \quad (7)$$

2. **(List decoding):** The decoder then constructs a list \mathcal{L} based on the prefix \mathbf{Y}_1^τ . More specifically, message m is put in the list if

$$|\{i \in [1 : \tau] : u_i(m) \neq Y_i, Y_i \neq \perp\}| < n^{3/4}. \quad (8)$$

That is, all codewords which are sufficiently close in Hamming distance (on the unerased bits) are put in the list.

3. **(List disambiguation):** The decoder then turns to the suffix $\mathbf{Y}_{[(\tau+1):n]}$. For a tuple (m_1, m_2, k_1, k_2) define the set of unerased bits that are in the k_1 -th part of m_1 and the k_2 -th part of m_2 :

$$V_{m_1, m_2, k_1, k_2} = \{i \in [(\tau + 1) : n] \cap S(m_1, k_1) \cap S(m_2, k_2) : Y_i \in \{0, 1\}\}. \quad (9)$$

For each pair $(m_1, m_2) \in \mathcal{L} \times \mathcal{L}$, the decoder first checks to see if there exists a (k_1, k_2) such that $k_1 \neq k_2$ and

$$|V_{m_1, m_2, k_1, k_2}| \geq \frac{\epsilon n}{4(K^2 - K)}. \quad (10)$$




Figure 2: Encoder: the message m is encoded first to $\mathbf{u}(m)$ and then noise \mathbf{Z} is added according to the subsets $\{S(m, k)\}$ and corresponding noise level probabilities $\{q_k\}$. The subsets are represented by different shades of grey. The resulting codeword \mathbf{X} is corrupted by Calvin who may use at most pn erasures. The received word is \mathbf{Y} .

If no such pair (k_1, k_2) , $k_1 \neq k_2$, exists then the decoder declares a decoding error. If such a pair exists the decoder takes the first such pair (lexicographically ordered) over all $K^2 - K$ such pairs, which we denote by V_{m_1, m_2} .

We adopt a simplified maximum likelihood decoding rule. Partition the set of indices into positions where $\mathbf{u}(m_1)$ and $\mathbf{u}(m_2)$ agree or disagree:

$$V_0 = \{i \in V_{m_1, m_2} : u_i(m_1) = u_i(m_2)\} \quad (11)$$

$$V_1 = \{i \in V_{m_1, m_2} : u_i(m_1) \neq u_i(m_2)\}. \quad (12)$$

Set V to be the larger of the two sets so that $|V| \geq |V_{m_1, m_2}|/2$. We apply the maximum likely decoder to V . Let $\alpha(m) = d_H(\mathbf{Y}_V, \mathbf{u}_V(m))$. We say m_1 beats m_2 if

$$\frac{q_{k_1}^{\alpha(m_1)}(1 - q_{k_1})^{|V| - \alpha(m_1)}}{q_{k_2}^{\alpha(m_2)}(1 - q_{k_2})^{|V| - \alpha(m_2)}} > 1, \quad (13)$$

otherwise we say m_2 beats m_1 .

4. If there exists a message \hat{m} in the list \mathcal{L} that beats all other elements of the list (a ‘‘Condorcet winner’’) then output that message \hat{m} , else it declares an error.

3.2 Analysis

In the analysis we follow the usual recipe: we show that for sufficiently large n , with high probability, a randomly constructed code will have P_{err} that vanishes as $n \rightarrow \infty$, thereby showing that such a code exists. Recall that in our code construction, we choose the ‘pure codewords’ $\mathbf{u}(m)$ independently and uniformly at random from $\{0, 1\}^n$ (i.e. they are i.i.d. $\text{Bernoulli}(1/2)$); and we generate the partitions $\{S(m, k)\}_k$ for each m by binning the indices in $[n]$ into K bins uniformly at random. The codebook consists of both the ‘pure codewords’ as well as the partitions for each message.

Let the random variable representing this codebook be denoted by \mathcal{C} . We will prove that the codebook has nice properties with a probability that is super-exponentially close to 1.

Fix any $\epsilon > 0$ and recall $R = 1 - p - \epsilon$. We prove a sequence of lemmas to prove we can achieve rate R .

Lemma 1. *With probability at least $1 - \exp\left(-\frac{\sqrt{n}}{2}\right)$ over encoder’s random noise \mathbf{Z} , the Hamming weight of \mathbf{Z} is at most $n^{3/4}$.*

Proof. Since each $\mathbb{P}(Z_j = 1) \leq q_K$ for all j , the probability is upper bounded by the probability that n i.i.d. variables $A_j \sim \text{Bernoulli}(q_K)$ have Hamming weight greater than $n^{3/4}$. The expected weight of \mathbf{A} is $q_K n = (2^{(\log(n)/4) - 1} n^{-1/2}) n = \frac{n^{3/4}}{2}$. Therefore by Hoeffding’s inequality,

$$\mathbb{P}\left(\sum_j Z_j > n^{3/4}\right) \leq \exp\left(-\frac{\sqrt{n}}{2}\right). \quad (14)$$

□




Figure 3: An example demonstrating coherence over \mathcal{T} for two base codewords $\mathbf{u}(m)$ and $\mathbf{u}(m')$. Let $\mathcal{T} = [25]$, *i.e.*, it comprises of the first 25 locations of each codeword, and let there be $K = 5$ noise-levels q_1, \dots, q_5 for each codeword, in the sets of locations $S(m, 1), \dots, S(m, 5)$ and $S(m', 1), \dots, S(m', 5)$ respectively. The expected size of each $V(m, m', k, k')$ is therefore $|\mathcal{T}|/K^2 = 1$. It can be verified that the largest size of $V(m, m', k, k')$ is 2 (only for $k = 3$ – all other sets of size at least 2 have $k \neq k'$), and hence $\mathbf{u}(m)$ and $\mathbf{u}(m')$ are at most 2-coherent over \mathcal{T} . Further, there are exactly 4 locations in which $\mathbf{u}(m)$ and $\mathbf{u}(m')$ are coherent (the 8th, 12th, 16th and 18th locations, as highlighted in this figure), so the remaining 21 decoherent locations are potentially usable by the decoder Bob, to disambiguate between m and m' . In this example, $V(m, m', 1, 2)$ comprising of the two locations $\{1, 20\}$ is a possible choice for the disambiguation set, being of “reasonable size”, and being the lexicographically first set with $k \neq k'$.

Lemma 2. *The length τ given in (7) of the prefix \mathbf{Y}_1^τ is at most $(1 - \epsilon/2)n$ and the suffix $\mathbf{Y}_{\tau+1}^n$ has at least $n\epsilon/2$ unerased bits.*

Proof. Since $R = 1 - p - \epsilon$ and the adversary can erase at most pn locations, the number of unerased bits in $\mathbf{Y}_1^{(1-\epsilon/2)n}$ is at least $(1 - p - \epsilon/2)n = n(R + \epsilon/2)$, as required by the definition of τ in (7). Let λ be the number of unerased bits in \mathbf{Y}_1^τ . By our definitions it holds that $\lambda = (1 - p - \epsilon/2)n$. Thus the number of erased bits in \mathbf{Y}_1^τ is $\tau - \lambda$. Implying at most $pn - \tau + \lambda$ erased bits in $\mathbf{Y}_{\tau+1}^n$, which finally implies at least

$$(n - \tau) - (pn - \tau + \lambda) = n - pn - (1 - p - \epsilon/2)n = \epsilon n/2 \quad (15)$$

unerased bits in $\mathbf{Y}_{\tau+1}^n$. \square

We now define a few useful properties of our random code. The decoder will have difficulty resolving the difference between codewords if they share very similar partitions $\{S(m, k)\}$. For two messages (m, m') , set $\mathcal{T} \subseteq [n]$, the expected number of common locations over \mathcal{C} is

$$\mathbb{E}_{\mathcal{C}} \left[\sum_{k=1}^K |S(m, k) \cap S(m', k) \cap \mathcal{T}| \right] = \sum_{k=1}^K \frac{|\mathcal{T}|}{K^2} = \frac{|\mathcal{T}|}{K}. \quad (16)$$

For a $\mathcal{T} \subset [n]$, call a pair of base codewords $(\mathbf{u}(m), \mathbf{u}(m'))$ η_1 -coherent over \mathcal{T} if

$$\sum_{k=1}^K |S(m, k) \cap S(m', k) \cap \mathcal{T}| \leq \frac{|\mathcal{T}|}{K} (1 + \eta_1) \quad (17)$$

That is, the number of locations in \mathcal{T} in which both $\mathbf{u}(m)$ and $\mathbf{u}(m')$ have the same noise levels is at most an $(1 + \eta_1)$ factor greater than the expected number of such locations. We call a codebook \mathcal{C} is (η_1, η_2) -coherent if for each pair of messages (m, m') and each \mathcal{T} of size at least $\eta_2 n$, the pair of base codewords $(\mathbf{u}(m), \mathbf{u}(m'))$ are η_1 -coherent over \mathcal{T} .

Let \mathcal{T}' be an ordered subset of $[n]$, and denote its i th entry by $(\mathcal{T}')_i$. Define the *restriction of a base codeword $\mathbf{u}(m)$ to \mathcal{T}'* , $\mathbf{u}_{\mathcal{T}'}(m)$, as the length- $|\mathcal{T}'|$ binary vector whose i -th entry equals the $(\mathcal{T}')_i$ th entry of $\mathbf{u}(m)$. Define the *restriction of a codebook \mathcal{C} to \mathcal{T}'* , denoted $\mathcal{C}_{\mathcal{T}'}$, is analogously defined as the codebook (with possible repetitions) generated by restricting each base codeword $\mathbf{u}(m) \in \mathcal{C}$ to \mathcal{T}' . We call codebook \mathcal{C} (w_u, w_e, s) -list-decodable if for each set $\mathcal{T}' \subset [n]$ of size at least w_u (of unerased bits), the restricted codebook $\mathcal{C}_{\mathcal{T}'}$ is “decodable against weight w_e errors to a list of size at most s ”. More precisely, for the (unrestricted) codebook \mathcal{C} , for any set $\mathcal{T}' \subset [n]$ of size at most w_u , any Hamming ball in $\{0, 1\}^{w_u}$ of radius at most w_e contains less than s codewords restricted to \mathcal{T}' .

Lemma 3. *For any sufficiently small $\epsilon > 0$, there exists sufficiently large N_ϵ , such that for all $n > N_\epsilon$, with probability at least*

$$1 - 2^{-\frac{\epsilon^2}{4} n \log \log(n)} \quad (18)$$

over the design of codebook \mathcal{C} , the following two properties hold:

1. The codebook \mathcal{C} is at most $(K/2 - 1, \epsilon/2)$ -coherent.
2. The codebook \mathcal{C} is $(n(1 - p - \epsilon/2), n^{3/4}, (\log \log(n))\epsilon/2)$ -list-decodable.

Proof. We first prove that with high probability \mathcal{C} is at most $(K/2 - 1, \epsilon/2)$ -coherent. Since $\eta_1 = (K/2 - 1)$, we must show that for any pair of base codewords $(\mathbf{u}(m), \mathbf{u}(m'))$

$$\sum_{k=1}^K |S(m, k) \cap S(m', k) \cap \mathcal{T}| \leq \frac{|\mathcal{T}|}{2}. \quad (19)$$

Recall that the sets $\{S(m, k)\}_{k=1}^K$ partition $[n]$. We will calculate the probability over the randomly selected partitions $\{S(m, k)\}$.

Fix any set $\mathcal{T} \subset [n]$ with size at least $\epsilon n/2$. The probability that the code construction generates $\frac{|\mathcal{T}|}{2}$ or more positions in which m and m' select the same q_k can be written as follows:

$$\sum_{i=\frac{|\mathcal{T}|}{2}}^n \binom{n}{i} \left(\frac{1}{K}\right)^i \left(1 - \frac{1}{K}\right)^{n-i} < \sum_{i=\frac{|\mathcal{T}|}{2}}^n \binom{n}{i} \left(\frac{1}{K}\right)^i \quad (20)$$

$$< \sum_{i=\frac{|\mathcal{T}|}{2}}^n 2^n \left(\frac{1}{K}\right)^i \quad (21)$$

$$< n 2^n \left(\frac{1}{K}\right)^{\frac{|\mathcal{T}|}{2}} \quad (22)$$

$$< n 2^n \left(\frac{4}{\log(n)}\right)^{\frac{\epsilon n}{4}} \quad (23)$$

$$= 2^{-\frac{\epsilon}{4} n \log \log(n) + (1 + \epsilon/2)n + \log(n)}, \quad (24)$$

where the last inequality follows from the setting of K as $\log(n)/4$ and the size of \mathcal{T} being at least $n\epsilon/2$. Taking a union bound over all pairs of base codewords (there are strictly less than 2^{2n} such pairs, since the rate of the code is less than 1) and all possible sets \mathcal{T} (there are strictly less than 2^n such sets) shows that the probability that a code is *not* at most $(K/2 - 1, \epsilon/2)$ -coherent is at most $2^{-\frac{\epsilon}{4}n \log \log(n) + (4+\epsilon/2)n + \log(n)}$.

We now prove that with high probability \mathcal{C} is appropriately list-decodable. This is broadly similar to classical derivations of list-decoding bounds, but due to the specific combination of error/erasure decoding required in this proof (with asymptotically vanishing fraction of errors but constant fraction of erasures) we re-derive a proof here. Since each base codeword $\mathbf{u}(m)$ in the codebook \mathcal{C} is generated uniformly at random from $\{0,1\}^n$, the same is true for codewords restricted to \mathcal{T}' (for any \mathcal{T}'). Therefore for all sufficiently large n , the probability that a codeword in $\mathcal{C}_{\mathcal{T}'}$ falls in any fixed Hamming ball of radius $n^{3/4}$ in $\{0,1\}^{n(1-p-\epsilon/2)}$ is

$$\frac{\binom{n(1-p-\epsilon/2)}{n^{3/4}}}{2^{n(1-p-\epsilon/2)}} = 2^{-n(1-p-\epsilon/2) + n^{3/4} \log(n^{1/4}) + \mathcal{O}(n^{1/4})} < 2^{-n(1-p-2\epsilon/3)} \quad (25)$$

where the equality follows from Stirling's approximation.

Let $\nu = 2^{-n(1-p-2\epsilon/3)}$. The probability (over the design of codebook \mathcal{C}) then that the Hamming ball contains at least $(\log \log(n))\eta_3$ codewords restricted to \mathcal{T}' is at most¹,

$$\sum_{i=(\log \log(n))\epsilon/2}^{2^{nR}} \binom{2^{nR}}{i} q^i (1-q)^{2^{nR}-i} < \sum_{i=(\log \log(n))\epsilon/2}^{2^{nR}} \binom{2^{nR}}{i} \nu^i \quad (26)$$

$$< \sum_{i=(\log \log(n))\epsilon/2}^{2^{nR}} 2^{nRi} \nu^i \quad (27)$$

$$= \sum_{i=(\log \log(n))\epsilon/2}^{2^{nR}} 2^{n(1-p-\epsilon)i} \left(2^{-n(1-p-2\epsilon/3)}\right)^i \quad (28)$$

$$< 2^{nR} 2^{-\frac{\epsilon^2}{3}n \log \log(n)} \quad (29)$$

$$< 2^{-\frac{\epsilon^2}{3}n \log \log(n) + n}. \quad (30)$$

Taking a union bound over all $2^{n(1-p-\epsilon/2)} < 2^n$ Hamming balls and all $\binom{n}{n(1-p-\epsilon/2)} < 2^n$ sets \mathcal{T}' implies that the probability (over design of \mathcal{C}) that there exists a set \mathcal{T}' for which there exists a Hamming ball with at least $(\log \log(n))\epsilon/2$ codewords restricted to \mathcal{T}' is at most $2^{-\frac{\epsilon^2}{3}n \log \log(n) + 3n}$.

Hence the probability that at least one of the two properties (approximate decoherence, and list-decodability) required do not hold for the codebook \mathcal{C} is at most $2^{-\frac{\epsilon}{4}n \log \log(n) + (4+\epsilon/2)n + \log(n)} + 2^{-\frac{\epsilon^2}{3}n \log \log(n) + 3n}$, which is at most $2^{-\frac{\epsilon^2}{4}n \log \log(n)}$ for all sufficiently small ϵ and sufficiently large n . \square

Lemma 4. *With probability at least*

$$1 - 2^{-\frac{\epsilon^2}{4}n \log \log(n)} \quad (31)$$

over the design of codebook \mathcal{C} , for any adversarial erasure pattern \mathbf{e} ,

¹Note that the expected number of codewords in the Hamming ball is no more than $q2^{nR}$, which equals $2^{-n\epsilon/3}$.

1. For every message m , the size of list \mathcal{L} in Equation (8) is at most $\log \log(n)\epsilon/2$ if $\text{wt}_H(\mathbf{Z}) \leq n^{3/4}$, and
2. For every pair of messages (m_1, m_2) , there exists a pair (k_1, k_2) , with $k_1 \neq k_2$, satisfying Equation (10).

Proof. 1. By Lemma 2, the prefix-length τ is at most $(1 - \epsilon/2)n$. Hence using part 2 of Lemma 3 gives us the required bound on list-decodability.²

2. We set \mathcal{T} to the indices corresponding to unerased bits in $[\tau + 1 : n]$ (which is of size at least $n\epsilon/2$ by Lemma 2). Part 3 of Lemma 1 shows that with probability at least $1 - 2^{-\frac{\epsilon^2}{4}n \log \log(n)}$, the size of the set $\bigcup_{k=1}^K S(m_1, k) \cap S(m_2, k) \cap \mathcal{T}$ is at most $\frac{|\mathcal{T}|}{2}$. Hence for any (m_1, m_2) the set

$$\left| \bigcup_{(k, k') \in [K] \times [K], k \neq k'} S(m_1, k) \cap S(m_2, k') \cap \mathcal{T} \right| \geq \frac{|\mathcal{T}|}{2} \geq \frac{n\epsilon}{4}. \quad (32)$$

But there are at most $K^2 - K$ values for the pair $(k, k') \in [K] \times [K]$ such that $k \neq k'$. Hence for at least one such pair, the size of $S(m_1, k) \cap S(m_2, k') \cap \mathcal{T}$ is at least $\frac{\epsilon n}{4(K^2 - K)}$, as required by (10). □

Lemma 5. *For a code satisfying the two conditions of Lemma 4, there exists a constant c such that for sufficiently large n , with probability at least $1 - \exp(-\epsilon\beta n^{1/2}/2 \log^2 n)$ over the encoder noise \mathbf{Z} , the decoder outputs the transmitted message m .*

Proof. Suppose m was transmitted and consider the test in the decoding rule for $m_1 = m$ and $m_2 = m' \neq m$. Let $q = q_{k_1}$ and $q' = q_{k_2}$. Let ζ denote the fraction of 1's in \mathbf{Z}_V (i.e. its type). If $V = V_0$ the decoder is a the maximum likelihood detector with $|V|$ observations between hypotheses $Z_i \sim \text{Bernoulli}(q)$ and $Z_i \sim \text{Bernoulli}(q')$. Message m beats m' if [7, (11.194)]:

$$D(\zeta \| q') - D(\zeta \| q) = \zeta \log \frac{q}{q'} + (1 - \zeta) \log \frac{1 - q}{1 - q'} \quad (33)$$

$$> 0. \quad (34)$$

If $V = V_1$ it is between $Z_i \sim \text{Bernoulli}(q)$ and $Z_i \sim \text{Bernoulli}(1 - q')$, so m beats m' if

$$D(\zeta \| 1 - q') - D(\zeta \| q) = \zeta \log \frac{q}{1 - q'} + (1 - \zeta) \log \frac{1 - q}{q'} \quad (35)$$

$$> 0. \quad (36)$$

In both cases we can solve for the ζ^* at the threshold (where the left side equals 0). By Sanov's Theorem [7, Theorem 11.4.1], the probability of error is

$$\mathbb{P}(m' \text{ beats } m) \leq |V| \exp(-|V|D(\zeta^* \| q)). \quad (37)$$

²In fact Lemma 3.2 provides stronger guarantees than are required in this proof. For one, it shows list-decodability for *any* \mathcal{T} of appropriate size, whereas the decoder only ever decodes using Y_1^τ . Furthermore, part 2 of Lemma 3 guarantees that *any* Hamming ball of appropriate radius does not correspond to too many messages, rather than just those Hamming balls centred at sub-vectors of Y_1^τ . Neither of these relaxations asymptotically worsens the parameters obtainable in this proof, but they have the advantage of significantly simplifying presentation.

Thus we must lower bound the divergence in both cases. Since $q, q' \ll 1/2$ it is clear that the case $V = V_0$ will have a smaller upper bound, so we focus on that case. For $V = V_0$ the error is largest when the hypotheses are closest, so $|k_1 - k_2| = 1$.

We first prove a useful lower bound on divergences. Using Taylor expansion, for $r \in (0, 1)$ and $\lambda > 0$ such that $\lambda r < 1$,

$$D(\lambda r \| r) = \lambda r \ln \lambda + (1 - \lambda r) \ln \frac{1 - \lambda r}{1 - r} \quad (38)$$

$$= \lambda r \ln \lambda + (1 - \lambda r) \left(\sum_{j=1}^{\infty} \frac{r^j}{j} - \sum_{j=1}^{\infty} \frac{\lambda^j r^j}{j} \right) \quad (39)$$

$$= \lambda r \ln \lambda + (1 - \lambda r)r + (1 - \lambda r) \sum_{j=2}^{\infty} \frac{r^j}{j} - \lambda r + \sum_{j=2}^{\infty} \left(\frac{1}{j-1} - \frac{1}{j} \right) \lambda^j r^j \quad (40)$$

$$> r(\lambda \ln \lambda - \lambda + 1) - \lambda r^2. \quad (41)$$

Now, $\lambda \ln \lambda - \lambda + 1 = 0$ at $\lambda = 1$ and

$$\frac{d}{d\lambda}(\lambda \ln \lambda - \lambda + 1) = \ln \lambda \quad (42)$$

so the coefficient of r is strictly positive for all $\lambda \neq 0, 1$. Thus for sufficiently small r , for any $\lambda > 0$, $\lambda \neq 1$ there exists a $\beta > 0$ such that $D(\lambda r \| r) \geq \beta r$.

Now we will apply this to our divergence for the threshold. We either have $q' = q/2 < \zeta^* < q$ or $q < \zeta^* < q' = 2q$, which means $r < \zeta^* < 2r$ for $r = q/2$ or q , and $D(\zeta^* \| r) = D(\zeta^* \| 2r)$. Therefore either $|\zeta^* - r| > r/2$ or $|2r - \zeta^*| > r/2$, which means that $D(\zeta^* \| r) > D(3r/2 \| r)$ or $D(\zeta^* \| 2r) > D(3r/2 \| 2r)$. In either case, the previous argument shows that there exists a $\beta > 0$ such that $D(\zeta^* \| r) \geq \beta r$. Therefore

$$D(\zeta^* \| q) \geq \beta q \geq \beta n^{-1/2}. \quad (43)$$

Let \mathcal{E}^c be the event that the conditions in Lemmas 1, 2, 3, and 4 hold. Taking a union bound over all messages m' in the list, we use the fact that with

$$\mathbb{P}(\text{any } m' \neq m \text{ beats } m | \mathcal{E}^c = c) \leq |\mathcal{L}| |V| \exp(-|V| D(\zeta^* \| q)) \quad (44)$$

$$= (\log \log n) \frac{\epsilon n}{\log^2 n} \exp\left(-\epsilon \beta n^{1/2} / \log^2 n\right) \quad (45)$$

$$\leq \exp\left(-\epsilon \beta n^{1/2} / 2 \log^2 n\right). \quad (46)$$

□

These lemmas together imply Theorem 1 as argued below.

Proof of Theorem 1: The converse follows by considering an adversary that erases the first pn bits. Fix $\epsilon > 0$ and set $R = 1 - p - \epsilon$. Fix any erasure pattern \mathbf{e} . By Lemma 4, with probability at least (31) the list \mathcal{L} contains at most $\log \log(n) \epsilon / 2$ codewords if $\text{wt}_H(\mathbf{Z}) < n^{3/4}$ and for each (m, m') in the list there exists a set $V(m, m', k, k')$ of size at least $\frac{\epsilon n}{4(K^2 - K)} = O(\epsilon n / \log^2 n)$. We union bound over all erasure patterns \mathbf{e} and messages m to show that with high probability the code construction satisfies these conditions. To complete the proof, note that by Lemma 1 the

weight of \mathbf{Z} is such that the list size is at most $\log \log(n)\epsilon/2$ with probability $1 - \exp(-n^{1/2}/2)$. Thus from Lemma 5 decoding succeeds with probability $1 - \exp(-c\epsilon n^{1/2} \log \log n / \log^3 n)$ for some $c > 0$. Therefore the probability of error goes to 0 as $n \rightarrow \infty$, showing that R is achievable. This completes the proof of Theorem 1.

4 Deterministic codes: Proof of Theorem 2

In this section we show that the stochastic nature of our code design is essential. Specifically, we show that any series of $(n, 2^{nR})$ *deterministic* codes (Φ_n, ψ_n) (i.e., for which $\Phi_n : [2^{nR}] \rightarrow \mathcal{X}^n$ depends only on $m \in [2^{nR}]$) that allow communication over our channel model with average probability of error ε_n tending to zero must satisfy $R \leq 1 - 2p$. An example illustrating our proof appears at the end of the section (see Figure 4)

We show, by presenting an adversarial strategy, that for any constant $\delta > 0$ and sufficiently large values of n , any deterministic code (Φ_n, ψ_n) with $R = 1 - 2p + \delta$ will have average error $\varepsilon_n = \Omega(1)$ (where ε_n does not depend on n but will depend on δ). The adversarial strategy follows the “wait and push strategy” (used in [19, 20] for the causal binary bit-flip channel and in [1] for the erasure case) in which the adversary “waits” a certain amount of time without performing any action, and then based on the information the adversary has seen so far “pushes” (i.e., corrupts) the transmitted codeword in a malicious manner causing a decoding error with some probability.

For a given message m and time parameter ℓ , let $\Phi_\ell(m)$ be the set of messages that have corresponding codewords that agree with $\Phi(m)$ on the first ℓ entries. The set $\Phi_\ell(m)$ plays an important role in our analysis and will be referred to as the “ ℓ -consistency” set. Notice that Calvin *cannot* construct $\Phi_\ell(m)$ after ℓ bits of $\Phi(m)$ have been transmitted (as, due to the delay, he has no knowledge of the ℓ 'th bit in $\Phi(m)$). However, as the delay of Calvin is only 1-bit, at each time step ℓ , Calvin will can construct two *potential* consistency sets. The set $\Phi_\ell^0(m)$ corresponding to the case that the ℓ 'th bit transmitted is 0 and one set $\Phi_\ell^1(m)$ corresponding to the case that the bit is 1. It holds that $\Phi_\ell^0(m) \cup \Phi_\ell^1(m) = \Phi_{\ell-1}(m)$.

We start by defining (and analyzing) the “wait” phase of Calvin. We will then turn to discussing the push phase.

4.1 “Wait” phase

In the wait phase Calvin proceeds as follows:

1. (*Wait-1*): Calvin starts by waiting until $(R - \delta)n = (1 - 2p + \delta)n + 1$ bits of the transmitted codeword are sent.
2. For each value of $\ell > (1 - 2p + \delta)n$, on transmission of the ℓ 'th bit of the transmitted codeword, Calvin constructs the sets $\Phi_\ell^0(m)$ and $\Phi_\ell^1(m)$. Let $A_\ell = \max(|\Phi_\ell^0(m)|, |\Phi_\ell^1(m)|)$ and $a_\ell = \min(|\Phi_\ell^0(m)|, |\Phi_\ell^1(m)|)$. Clearly $A_\ell \geq a_\ell$. In addition, $A_\ell + a_\ell$ is exactly $|\Phi_{\ell-1}(m)|$ and we will show shortly that with high probability over messages m it holds that for $\ell = (1 - 2p + \delta)n$ the size $A_\ell + a_\ell$ is at least $2^{\Theta(n)}$. Based on the value of $A_\ell + a_\ell$ Calvin decides to either continue waiting or to move on to the push phase. Specifically:
 - (*Wait-2*): If $A_\ell + a_\ell$ is greater than $\delta'n$, Calvin does nothing and waits for the next bit to be transmitted. Here $\delta' = \delta/4$.

- (*Attack*): If $A_\ell + a_\ell$ is less than $\delta'n$ but at least as large as $\frac{c}{\delta}$ Calvin, sets the “transition time” ℓ^* to equal the current value of ℓ , stops the “wait” phase, and moves on to the “push” phase to be discussed below in Section 4.2 in detail.
- (*Error*): If $A_\ell + a_\ell$ is less than $\frac{c}{\delta}$, set the “transition time” ℓ^* to equal the current value of ℓ and declare an error of Type 1.

By our definitions, Calvin either declares an error or will move on to the push phase at some point in time ℓ^* . For the latter we say that the transition to the push phase is successful for message m , namely the size $A_{\ell^*} + a_{\ell^*} < \delta'n$ is at least $\frac{c}{\delta}$ for a sufficiently large constant c to be determined shortly. Otherwise we say that the transition has failed (there is an error of Type 1). We now show that with some constant probability over messages m , the transition to the push phase is successful (no error of Type 1).

Lemma 6. *Let c be a sufficiently large constant to be determined shortly. Let n be sufficiently large. Let ℓ^* be the first point in time for which $A_{\ell^*} + a_{\ell^*} < \delta'n$. With probability at least $2^{-8c/\delta^2}$ over messages m , it holds that $A_{\ell^*} + a_{\ell^*}$ is of size at least $\frac{c}{\delta}$.*

Proof. We first note that in [21] it is shown, using the pigeonhole principle, that the probability over messages m that for $\ell = (1 - 2p + \delta)n$ the size of $\Phi_\ell(m)$ (and thus $A_\ell + a_\ell$) is at least $2^{\delta n/2}$ is at least $1 - 2^{-\delta n/2}$. Let E_1 be the event that Alice chooses a message m for which the corresponding consistency set $\Phi_\ell(m)$ is of size at least $2^{\delta n/2}$.

We now address the probability, given E_1 that the transition of Calvin to the push phase has failed. This can happen for messages m only if $\Phi_{\ell^*-2}(m) = A_{\ell^*-1} + a_{\ell^*-1} \geq \delta'n$ and $\Phi_{\ell^*-1}(m) = A_{\ell^*} + a_{\ell^*}$ is of size less than $\frac{c}{\delta}$. Or in other words, failure happens only for messages m that at some point in time have consecutive consistency sets of sizes that *jump* from above $\delta'n$ to below $\frac{c}{\delta}$.

Consider a codeword chosen uniformly at random from the codebook of Alice (this corresponds to choosing a uniformly distributed message m). One may expose this codeword bit by bit according to the conditional probability given the choices made thus far. In such a process for time parameter ℓ , if $A_{\ell-1} + a_{\ell-1}$ is at least $\delta'n$ and the value of $A_\ell + a_\ell$ is less than $\frac{c}{\delta}$, there will be a failure for Calvin with probability

$$\frac{A_\ell + a_\ell}{A_{\ell-1} + a_{\ell-1}}. \quad (47)$$

Notice that $A_\ell + a_\ell$ is equal to either $A_{\ell-1}$ or $a_{\ell-1}$ by our exposure process. Moreover, for sufficiently large n , $A_\ell + a_\ell = a_{\ell-1}$ as otherwise $A_{\ell-1} = A_\ell + a_\ell < \frac{c}{\delta}$ which in turn implies that $A_{\ell-1} + a_{\ell-1} \leq \frac{2c}{\delta}$ in contradiction to $A_{\ell-1} + a_{\ell-1} \geq \delta'n$. This implies that in such cases, the conditional probability of error at time ℓ is

$$\frac{a_{\ell-1}}{A_{\ell-1} + a_{\ell-1}}, \quad (48)$$

or equivalently, in such cases the conditional probability that the exposure process does *not* induce a failed transition is

$$1 - \frac{a_{\ell-1}}{A_{\ell-1} + a_{\ell-1}}. \quad (49)$$

We conclude that the probability q over codewords (i.e., messages m) that the transition is successful for Calvin is

$$q = \prod_{\ell} \left(1 - \frac{a_{\ell-1}}{A_{\ell-1} + a_{\ell-1}} \right), \quad (50)$$

where the product is over ℓ for which (as specified above) $A_{\ell-1} + a_{\ell-1} \geq \delta' n$ and $1 \leq a_{\ell-1} \leq \frac{c}{\delta}$. As there can be at most n such values of ℓ we have that

$$q \geq \prod_{k=1}^n \left(1 - \frac{c}{\delta x_k} \right), \quad (51)$$

where x_ℓ is a strictly decreasing sequence of integers greater than $\delta' n$. It now holds that the setting for which our lower bound on q is minimum is that in which x_k are consecutive integers (in increasing order) starting from $x_n = \delta' n + 1$. I.e. $x_k = \delta' n + k$. We conclude that (for sufficiently large values of n) q is bounded from below by $e^{-\left(\frac{4c}{\delta^2}\right)^3}$:

$$q \geq \prod_{k=1}^n \left(1 - \frac{c}{\delta(\delta' n + k)} \right) \geq \prod_{k=1}^n \left(1 - \frac{c}{\delta \delta' n} \right) = \left(1 - \frac{4c}{\delta^2 n} \right)^n \geq e^{-\left(\frac{4c}{\delta^2}\right)}.$$

All in all, using the union bound with event E_1 , for sufficiently large n we have with probability at least $e^{-\left(\frac{4c}{\delta^2}\right)} - 2^{-\delta n/2} \geq 2^{-\left(\frac{8c}{\delta^2}\right)}$ that Calvin's transition to the push phase will result in a success. \square

4.2 “Push” phase

Calvin's corrupting algorithm now proceeds as follows.

1. Calvin chooses a “plausible transmission” \mathbf{X}' uniformly at random from $\Phi_{\ell^*}^0(m) \cup \Phi_{\ell^*}^1(m)$.
2. For each value of $\ell \geq \ell^*$, either $|\Phi_{\ell}^0(m)| \geq 1$ and $|\Phi_{\ell}^1(m)| \geq 1$ (Calvin has uncertainty about X_{ℓ} , since it is possible for X_{ℓ} to equal either 0 or 1), or $|\Phi_{\ell}^i(m)| = 0$ for some $i \in \{0, 1\}$ (Calvin is certain about X_{ℓ} , since all surviving codewords have $\mathbf{X}_{\ell} = 1 - i$). Calvin does the following:
 - (a) (*Calvin uncertain about \mathbf{X}_{ℓ}*): If $|\Phi_{\ell}^0(m)| \geq 1$ and $|\Phi_{\ell}^1(m)| \geq 1$, then Calvin erases \mathbf{X}_{ℓ} .
 - (b) (*Calvin certain about \mathbf{X}_{ℓ}*): If $|\Phi_{\ell}^i(m)| = 0$ for some i ,
 - i. (*Erasing disambiguating information*): If $\mathbf{X}'_{\ell} = i$ and hence $\mathbf{X}_{\ell} \neq \mathbf{X}'_{\ell}$, Calvin erases \mathbf{X}_{ℓ} .
 - ii. (*No action*): If $\mathbf{X}'_{\ell} = 1 - i$ and hence $\mathbf{X}_{\ell} = \mathbf{X}'_{\ell}$, Calvin does not erase \mathbf{X}_{ℓ} .

If Calvin can successfully continue the above process until the end without violating his total erasure budget, then clearly both the codewords \mathbf{X} and \mathbf{X}' are consistent with the vector received by Bob as all the indices where they differ are erased by Calvin. We will now argue that Calvin can indeed complete this process, that is, the total number of erasures required is $\leq pn$.

³A tighter analysis indicates a better lower bound of $\left(\frac{\delta}{8}\right)^{\frac{c}{\delta}}$ – due to the intricacy of this analysis we omit it here.

Erasures are introduced by Calvin in steps 2(a) and 2(b)ii. Let us consider the full binary tree of depth n where the edges are labeled by 0 and 1, and let us consider the code as its subtree. Here, each path from the root to a leaf represents the codeword that is composed of the bits labeling the branches along that path. From Calvin's perspective, at the beginning of the push phase, the encoder state is either of the two nodes representing the subsequences $\mathbf{X}^{\ell^*-1}0$ and $\mathbf{X}^{\ell^*-1}1$. All the paths via these two nodes represent the two sets of codewords $\Phi_{\ell^*}^0(m)$ and $\Phi_{\ell^*}^1(m)$ respectively. Since the total number of such paths is at most $\delta'n = \delta n/4$, there are at most $\delta n/4$ branchings in the subtree rooted at the node corresponding to $\Phi_{\ell^*-1}(m)$ (i.e., the subtree spanning the codewords in $\Phi_{\ell^*}^0(m)$ and $\Phi_{\ell^*}^1(m)$). This implies, that along any path in this subtree Calvin will encounter at most $\delta n/4$ branching nodes. In other words, Calvin will encounter step 2(a) at most $\delta n/4$ times. This upper bounds the total number of erasures due to step 2(a) by $\delta n/4$.

Counting the number of required erasures in step 2(b)ii can be done following similar analysis as in [21], using the Plotkin bound and Turan's theorem. We briefly reprise the analysis here. We consider the codebook of length $n' \leq (2p - \delta)n$ formed by the completions of \mathbf{X}^{ℓ^*-1} . Let us consider the graph with these codewords as nodes, and two codewords connected if their Hamming distance is at most $d = pn - \delta n/4$. By the Plotkin bound, any independent set in this graph has at most $4p/\delta$ nodes. This implies, by Turan's theorem, that the average degree Δ and the number of nodes $|V|$ satisfy

$$\frac{\Delta + 1}{|V|} \geq \frac{\delta}{4p}.$$

Implying that

$$\frac{\Delta}{|V|} \geq \frac{\delta}{8p}.$$

Thus the probability of two randomly chosen codewords being at a distance at most d is

$$\frac{|\mathcal{E}|}{|V|^2} = \frac{\Delta|V|}{2|V|^2} \geq \frac{\delta}{16p}.$$

So with this constant probability, Calvin's remaining erasure budget $pn - \delta n/4$ is sufficient to erase (in step 2(b)ii) all the positions where \mathbf{X} and \mathbf{X}' differ.

All in all, the success probability of Calvin is bounded by below by his success in the wait phase times that in the push phase which is a constant independent of n :

$$\frac{\delta}{16p} \cdot 2^{-\left(\frac{8c}{\delta^2}\right)} \geq 2^{-\left(\frac{16c}{\delta^2}\right)}. \quad (52)$$

Remark 1. We note that the proof of Theorem 2 does not hold for stochastic codes. While the wait phase may have an analogous analysis that fits the stochastic setting, the push phase breaks down. Specifically, a crucial part of the push phase is step 2(a) which erases any location in which there is some uncertainty on behalf of Calvin regarding the current symbol. In deterministic codes, step 2(a) may occur in only few locations, whereas in the stochastic setting the number of branchings in the subtree rooted at the node corresponding to Calvin's view so far may be large, and thus step 2(a) may be too costly. Indeed, in our code design for the achievability proof presented in Section 3, each and every location includes a branching point.




Figure 4: Illustration of proof for Theorem 2

4.3 Illustration of proof for Theorem 2

In Figure 4, we demonstrate a toy example showing an adversarial attack against a deterministic code of block-length $n = 9$, comprising of 5 codewords $\mathbf{X}, \mathbf{X}', \mathbf{X}'', \mathbf{X}'''$ and \mathbf{X}'''' , and hence the rate R of this code equals $\log_2(5)/9 \approx .258$. In this toy-example, $p = 4/9$ (so at most 4 erasures are possible over the length-9 transmission, though in this example only 3 bits are actually erased), and hence the claim of Theorem 2 is that for sufficiently large n , no rate asymptotically larger than $1 - 2p = 1/9$ is achievable, implying that “not too many more” than 2 messages can be reliably transmitted via a deterministic code. In particular, this example aims to show that for the specific code shown, the 5 messages corresponding to the 5 codewords chosen cannot be reliably transmitted. (To keep the example dimensions manageable, not all parameters in the example match those in our proofs – in particular, no suitable value of the “rate-excess parameter δ exists that matches those required by our proofs, for the “small” value of n chosen.) The zig-zag lines at the bottom of Figure 4 show the “code-tree”, the binary tree representing the 5 length-9 codewords as paths in an (incomplete) depth-9 binary tree – segments angled upwards indicate 0’s in that location, and segments angled downwards indicate 1’s in that location, and hence the five codewords are respectively $\mathbf{X} = 100101011$, $\mathbf{X}' = 100111010$, $\mathbf{X}'' = 1001010001$, $\mathbf{X}''' = 100001010$, and $\mathbf{X}'''' = 111101010$.

The codeword actually transmitted, \mathbf{X} , is shown as the black-shaded path in the code-tree. Note that all codewords in the code have the same first bit (1), and hence Calvin has to wait until $\ell = 2$ before he knows that the transmitted codeword is not \mathbf{X}'''' . In general, Calvin’s initial two phases are “wait” phases, in which he does not erase any bits. Specifically, Calvin is *always* in the “Wait-1” phase for *exactly* the first $(1 - 2p + \delta)n + 1$ bits (shown in this figure as the first 3 bits),

and then he continues waiting in the “Wait-2” phase until time ℓ^* (when the number of codewords consistent with his observations up to the time $\ell^* - 1$ is somewhere in the range $(c/\delta, \delta n)$, for some constant c specified in Theorem 2). In this example the Wait-2 phase is of length 1, since at time $\ell^* = 5$ Calvin observes the 4th transmitted bit $x_4 = 1$, and realizes that the transmitted codeword does not equal \mathbf{X}''' – hence his consistency set at this time (denoted $\Phi_4(m)$) shrinks to become $\{\mathbf{X}, \mathbf{X}', \mathbf{X}''\}$, and is of size 3. At this point, Calvin segues to the “Attack” phase. Specifically, he first chooses a random codeword from his consistency set $\Phi_4(m)$ (in this example \mathbf{X}' , denoted as the shared black-red path in the code-tree), and tries to confuse Bob between \mathbf{X} and \mathbf{X}' . Specifically, whenever Calvin sees a “branch-point”, *i.e.*, a location ℓ in which Alice may have transmitted either a 0 or a 1 (*i.e.*, in which there are codewords corresponding to both $\Phi_\ell^0(m)$ and $\Phi_\ell^1(m)$) he erases the corresponding bit (Step 2(a) in the push phase of Theorem 2). This happens in at most $\Phi_{\ell^*-1}(m) < \delta n$ locations and thereby denies Bob knowledge of the value of these bits of \mathbf{X}_m (in this example, there are branch-points at $\ell = 5$ and $\ell = 8$). Also, if there is no branch-point, but $\mathbf{X}_\ell \neq \mathbf{X}'_\ell$ (such a bit would enable Bob to disambiguate between \mathbf{X} and \mathbf{X}'), Calvin erases such bits as well (this happens at $\ell = 9$). At the end, both \mathbf{X} and \mathbf{X}' are equally likely from Bob’s perspective. Care is required to ensure that Calvin does not run out of erasures in the Attack phase – this is analyzed in Theorem 2 in detail.

5 Omniscient adversary: stochastic vs. deterministic encoding

Here we argue that for a bit-flipping adversary who can flip upto p fraction of bits in a codeword, the capacity under stochastic encoding is the same as that under deterministic encoding.

Suppose a rate r is achievable under stochastic encoding and average error probability. Let us suppose that there is a sequence of stochastic codes achieving average probability of error ϵ_n for length n , such that $\epsilon_n \rightarrow 0$. Let us now consider a fixed n , and let \mathcal{A}_m denote the set of vectors for which the decoder outputs the message m . Let $\mathcal{E} := \{\hat{M} \neq M\}$.

$$\mathcal{M}_n := \{m : \Pr\{\mathcal{E} | M = m\} < 1\}.$$

For each message $m \in \mathcal{M}_n$, there is a ‘good’ codeword $\mathbf{X}(m)$ such that $\Pr\{\mathbf{X}(m) | M = m\} > 0$, and the adversary does not have the power to move it outside \mathcal{A}_m . In other words, the ball of radius pn around $\mathbf{X}(m)$ is completely contained in \mathcal{A}_m . Let $\alpha_n := \Pr\{M \in \mathcal{M}_n\}$. Clearly,

$$\begin{aligned} \epsilon_n &= \alpha_n \Pr\{\mathcal{E} | M \in \mathcal{M}_n\} + (1 - \alpha_n) \Pr\{\mathcal{E} | M \notin \mathcal{M}_n\} \\ &\geq 0 + (1 - \alpha_n) \\ &\geq 1 - \alpha_n. \end{aligned}$$

Let $\mathcal{C}_n := \{\mathbf{X}(m) : m \in \mathcal{M}_n\}$ be a set of good codewords for messages in \mathcal{M}_n . We now argue that the sequence of deterministic codes \mathcal{C}_n with decoder decision regions $\mathcal{A}_m; m \in \mathcal{M}_n$ have zero error, and an asymptotic rate r . That the code has zero error probability follows because, for each

codeword $\mathbf{X}(m)$, the adversary does not have the power to move it outside \mathcal{A}_m .

$$\begin{aligned}
H(M) &\geq nr \\
\Rightarrow H(\alpha_n) + \alpha_n H(M|M \in \mathcal{M}_n) + (1 - \alpha_n) H(M|M \notin \mathcal{M}_n) &\geq nr \\
\Rightarrow \alpha_n H(M|M \in \mathcal{M}_n) &\geq nr - 1 - (1 - \alpha_n) H(M|M \in \mathcal{M}_n) \\
\Rightarrow \alpha_n H(M|M \in \mathcal{M}_n) &\geq nr - 1 - \epsilon_n H(M|M \in \mathcal{M}_n) \\
\Rightarrow \alpha_n H(M|M \in \mathcal{M}_n) &\geq nr - 1 - \epsilon_n nr \\
\Rightarrow \frac{1}{n} H(M|M \in \mathcal{M}_n) &\geq \frac{1}{\alpha_n} \left((1 - \epsilon_n)r - \frac{1}{n} \right) \\
\Rightarrow \frac{1}{n} \log_2 |\mathcal{M}_n| &\geq \frac{1}{\alpha_n} \left((1 - \epsilon_n)r - \frac{1}{n} \right)
\end{aligned}$$

Since $\epsilon_n \rightarrow 0$ and $\alpha_n \rightarrow 1$, the rate of \mathcal{C}_n converges to r .

References

- [1] Raef Bassily and Adam Smith. Causal Erasure Channels. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1844–1857, 2014.
- [2] Zitan Chen, Sidharth Jaggi, and Michael Langberg. A characterization of the capacity of online (causal) binary channels. In *Foundations of Computer Science*, 2015.
- [3] Claude E. Shannon and Warren Weaver. *The mathematical theory of communication (Urbana, IL)*. University of Illinois Press IL, 1949.
- [4] David Blackwell, Leo Breiman, and A. J. Thomasian. The capacities of certain channel classes under random coding. *The Annals of Mathematical Statistics*, pages 558–567, 1960.
- [5] Amos Lapidoth, Prakash Narayan, and others. Reliable communication under channel uncertainty. *IEEE Transactions on Information Theory*, 44(6):2148–2177, 1998.
- [6] Imre Csiszár and Prakash Narayan. The capacity of the arbitrarily varying channel revisited: Positivity, constraints. *IEEE Transactions on Information Theory*, 34(2):181–193, 1988.
- [7] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [8] R. Ahlswede and J. Wolfowitz. Correlated decoding for channels with arbitrarily varying channel probability functions. *Information and Control*, 14:457–473, 1969.
- [9] Anand D. Sarwate and Michael Gastpar. Rateless codes for AVC models. *IEEE Transactions on Information Theory*, 56(7):3105–3114, 2010.
- [10] E. N. Gilbert. A comparison of signalling alphabets. *Bell Systems Technical Journal*, 31:504–522, 1952.
- [11] R. R. Varshamov. Estimate of the number of signals in error correcting codes. *Dokl. Acad. Nauk*, 117:739–741, 1957.

- [12] Robert J. McEliece, Eugene R. Rodemich, Howard Rumsey Jr, and Lloyd R. Welch. New upper bounds on the rate of a code via the Delsarte-MacWilliams inequalities. *IEEE Transactions on Information Theory*, 23(2):157–166, 1977.
- [13] I. Csiszár and P. Narayan. Arbitrarily varying channels with constrained inputs and states. *IEEE Transactions on Information Theory*, 34(1):27–34, 1988.
- [14] R. Ahlswede. Elimination of correlation in random codes for arbitrarily varying channels. *Z. Wahrsch. Verw. Gebiete*, 33:159–175, March 1978.
- [15] Michael Langberg. Private codes or succinct random codes that are (almost) perfect. In *FOCS*, volume 4, pages 325–334, 2004.
- [16] Adam Smith. Scrambling adversarial errors using few random bits, optimal information reconciliation, and better private codes. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 395–404. Society for Industrial and Applied Mathematics, 2007.
- [17] Aaron D Wyner. The wire-tap channel. *Bell System Technical Journal*, The, 54(8):1355–1387, 1975.
- [18] B. K. Dey, S. Jaggi, and M. Langberg. Codes against online adversaries, Part I: Large alphabets. *IEEE Transactions on Information Theory*, 59(6):3304–3316, 2013.
- [19] Bikash Kumar Dey, Sidharth Jaggi, Michael Langberg, and Anand D. Sarwate. Improved upper bounds on the capacity of binary channels with causal adversaries. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pages 681–685. IEEE, 2012.
- [20] Bikash Kumar Dey, Sidharth Jaggi, Michael Langberg, and Anand D. Sarwate. Upper bounds on the capacity of binary channels with causal adversaries. *IEEE Transactions on Information Theory*, 59(6):3753–3763, 2013.
- [21] Michael Langberg, Sidharth Jaggi, and Bikash Kumar Dey. Binary causal-adversary channels. In *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, pages 2723–2727. IEEE, 2009.